

LA-UR-19-28599

Approved for public release; distribution is unlimited.

Title: Improved Time-Stepping Methods in Global to Regional Ocean Modeling

Author(s): Petersen, Mark Roger
Wolfram, Phillip Justin Jr.
Bishnu, Siddhartha

Intended for: LANL Institutional Computing Proposal

Issued: 2019-08-25

Disclaimer:

Los Alamos National Laboratory, an affirmative action/equal opportunity employer, is operated by Triad National Security, LLC for the National Nuclear Security Administration of U.S. Department of Energy under contract 89233218CNA000001. By approving this article, the publisher recognizes that the U.S. Government retains nonexclusive, royalty-free license to publish or reproduce the published form of this contribution, or to allow others to do so, for U.S. Government purposes. Los Alamos National Laboratory requests that the publisher identify this article as work performed under the auspices of the U.S. Department of Energy. Los Alamos National Laboratory strongly supports academic freedom and a researcher's right to publish; as an institution, however, the Laboratory does not endorse the viewpoint of a publication or guarantee its technical correctness.

Improved Time-Stepping Methods in Global to Regional Ocean Modeling

PI: Mark R. Petersen, CCS-2, mpetersen@lanl.gov, Z#204932
 Phillip J. Wolfram, T-3, pwolfram@lanl.gov, Z#285088
 Siddhartha Bishnu, CCS-2, siddhartha.bishnu@lanl.gov, Z#322556

Proposed moniker: ocean_time_step

Abstract

Improvements in time-stepping schemes can provide substantial improvements in the throughput of global ocean-climate models, but is challenging due to large variations in time scales in the dynamics and spatial scales in the global, variable resolution mesh. Two funded projects aim to speed up the Model for Prediction Across Scales-Ocean (MPAS-Ocean). The first will add a variable-timestepping formulation to MPAS-Ocean, which is a critical step towards running efficiently on global variable-resolution meshes with regional and coastal refinement. The second optimizes the barotropic (depth-integrated) mode, which is currently a performance bottleneck at high core counts. The proposed compute time will be used for extensive model validation and parameter studies for both methods.

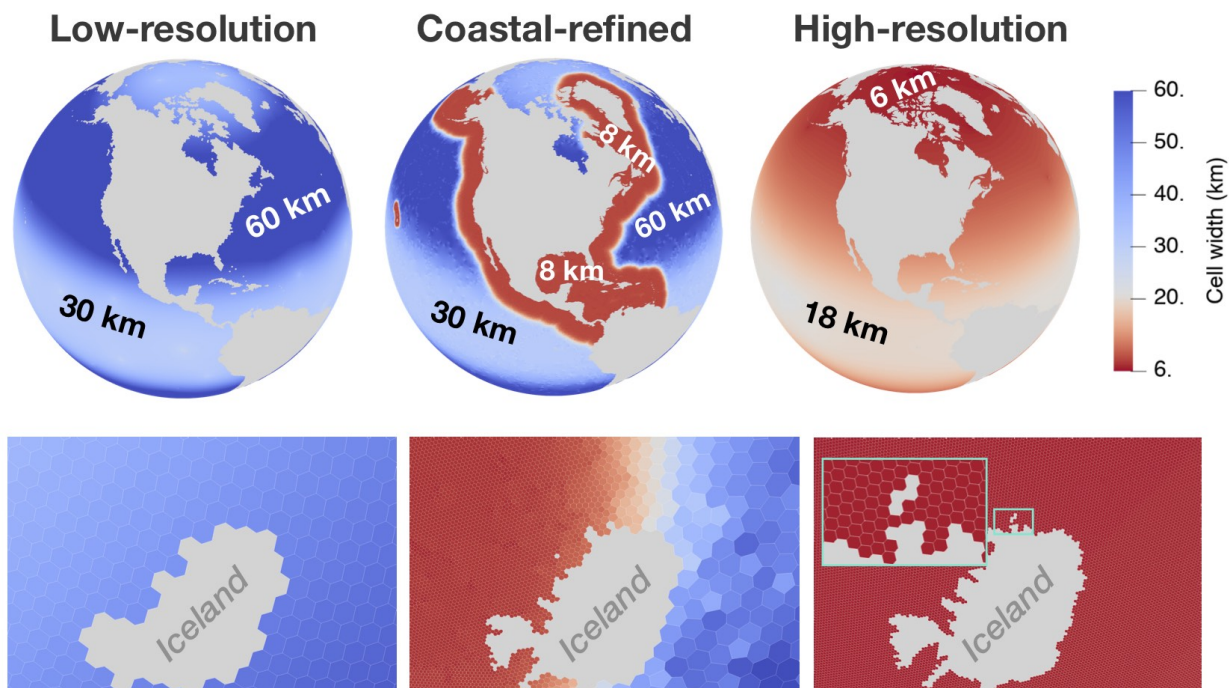


Figure 1: The MPAS-Ocean Coastal US-Plus 8 km (CUSP8) variable-resolution mesh uses the standard low-resolution mesh globally, but refines the grid cell size from 60 km to 8 km along the coast of North America. The polygonal Voronoi cells vary smoothly in size, as shown in the lower panels.

Technical Description: The Energy Exascale Earth System Model (E3SM) is the DOE’s new global coupled climate model, first released in 2019 [5, 2]. All components are based on variable-resolution meshes (Figure 1), so that E3SM offers unprecedented versatility in climate modeling applications focused on particular regions and processes (Figure 2). The ocean, sea ice, and land ice components are all developed at LANL, and are part of the Model for Prediction Across Scales (MPAS) framework. This includes a mesh specification, decomposition of variables across processors, parallel input and output specified in a run-time streams file, timers, and error handling. Finite volume operators were developed for Voronoi tessellations for

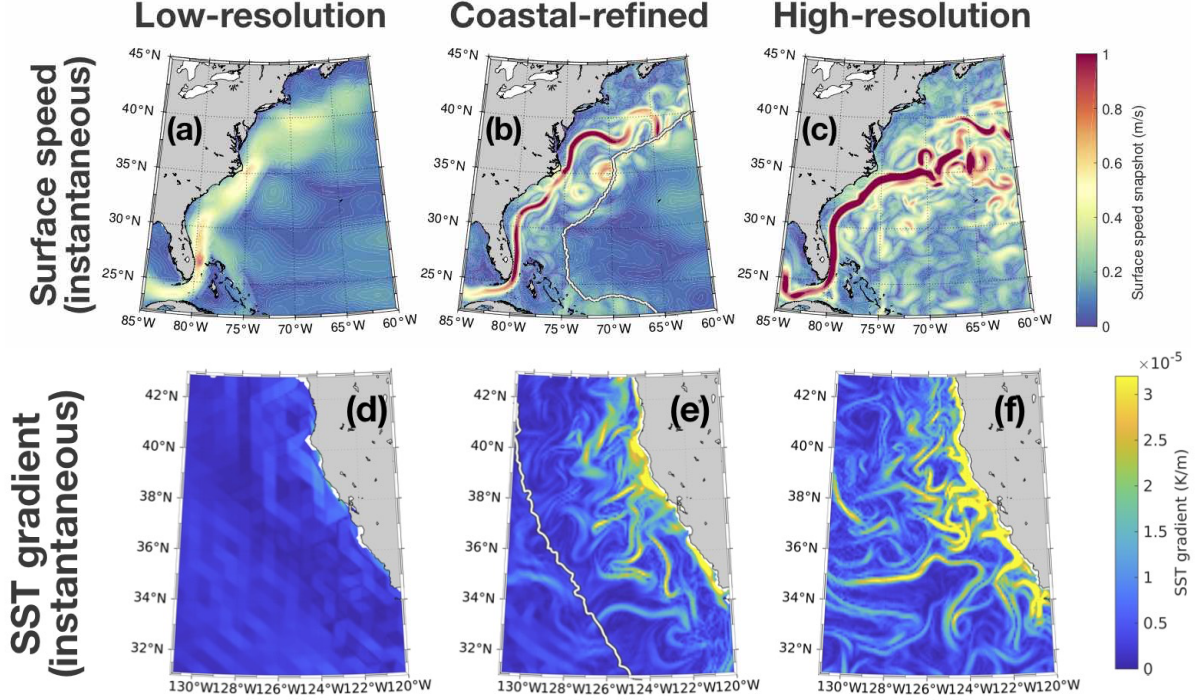


Figure 2: MPAS-Ocean simulation data for surface speed of the Gulf Stream (top) and sea surface temperature (SST) gradient for the California Current (bottom). The coastal-refined regions (middle panels, within white line) captures the fine-scale detail and variability of the global high-resolution simulation, while the low-resolution currents are weak and under-resolved.

the shallow water equations using mimetic methods to guarantee that mass, velocity and potential vorticity evolve in a consistent and compatible manner [8]. MPAS-Ocean solves prognostic equations for momentum, thickness (volume), and tracers using these operators [7] and can be run using both regular and unstructured meshes on Cartesian and spherical domains. MPAS-Ocean uses an Arbitrary Lagrangian-Eulerian (ALE) method for the vertical coordinate [5, 6]. It includes extensive in-situ analysis capabilities such as Lagrangian particles [9].

We propose to improve the efficiency of MPAS-Ocean with two new developments. The first is variable-resolution time-stepping, which will be funded by the new DOE effort on Integrated Coastal Modeling (LANL Earth System Model Development lead P. Wolfram). Currently, MPAS-Ocean uses a uniform timestep throughout the model domain. This is a hindrance to good performance, as the global time step is controlled by CFL condition of the smallest grid cell. It would be more efficient if low resolution regions could be run with longer time steps, based on the local CFL criterion. Variable-resolution time-stepping schemes partition the domain into several regions, with finer cells subcycling in time relative to larger cells. The challenge is to ensure mass conservation and wave propagation across these boundaries. University collaborator Lili Ju has developed an algorithm for this purpose that will be applied to MPAS-Ocean [3].

The second project is the optimization of the barotropic solver, a project funded by DOE BER through E3SM, with PI M. Petersen and graduate student S. Bishnu. Ocean models employ a split scheme between the 2D vertically-averaged barotropic component, and the remaining 3D baroclinic, because the barotropic includes fast surface gravity waves. MPAS-Ocean currently uses explicit subcycling for the barotropic component, and this project experiments with how the subcycled fields are filtered in time to produce the state at the next large time step. The best formulation could reduce the number of subcycles by 25-30%, which is a substantial performance gain because each subcycle requires an MPI communication for the halo update, and this causes a performance bottleneck at high core counts.




Figure 3: Performance versus resolution study on the LANL grizzly cluster, showing simulated years per wall clock day (SYPD). Black dotted lines show perfect scaling. The SYPD values for 1024 processors are: CUSP8: 2.0, CUSP12: 5.1, CUSP20: 11.7, CUSP30: 20.0, EC60to30 (low resolution): 32.5, RRS18to6 (high resolution): 0.38. All simulations use 80 layers, except the EC60to30, which is 60 layers. CUSP X is the Coastal US-Plus mesh shown in Figure 1, where X is the coastal grid cell size

Project Significance E3SM was initiated by DOE in 2013 to address the grand challenge of actionable predictions of earth system variability and change, with an emphasis on the most critical scientific questions facing the nation and DOE. Long-term science drivers include coastal flooding, sea level rise, water availability, and changes in storm frequency and extreme weather, all of which are critical to the energy sector. The project identified three major limitations that must be addressed to make progress on these goals: pushing the high-resolution frontier of Earth system modeling; bridging the gap in scales and processes; and ensemble modeling to quantify uncertainty. These goals require both higher fidelity modeling and higher computational throughput in the regions of interest. The strategy set forth for E3SM is a coupled climate model where all components are variable-resolution and optimized for high performance computing platforms.

The first set of E3SM simulations were based on meshes with modest variations in resolution, with grid cell size varying smoothly in latitude by a factor of two or three globally [4, 2]. We are now pushing towards more strongly varied meshes, with a factor of ten or more, but are limited by the performance issues described above.

HPC Parallel Computing Details: MPAS-Ocean development began at LANL in 2010 by PI Petersen and others, and was designed from the beginning for high performance computing platforms using both MPI and OpenMP. Low-resolution meshes scale well to thousands of cores (Figure 3), and high-resolution scales well to 50 thousand cores (tested on edison at NERSC). LANL IC has been an integral resource for MPAS component development. We use standard libraries including netcdf, pnetcdf, and parallel-io (PIO). All i/o is parallelized to take advantage of the distributed file system. MPAS-Ocean and MPAS-sea ice have extensive tools for mesh and initial condition generation [1], nightly regression testing, post-processing analysis and visualization. E3SM¹ and MPAS-Ocean² are open development projects, to facilitate collaboration with DOE and University partners.

E3SM is a large collaborative effort amongst 85 DOE employees across eight laboratories, with 15 staff members at LANL as well as post-docs and students. E3SM has large allocations at NERSC, ALCF, and OLCF for full coupled climate simulations. We request allocations on LANL IC for smaller development

¹<https://e3sm.org> and <https://github.com/E3SM-Project/E3SM>

²<https://github.com/MPAS-Dev/MPAS-Model>

		hor cells	vertical	# cores	measured	number	number	total	Million	storage	scratch	archived
resolution	name	thousands	layers	for timing	SYPD grizzly	sims	yrs/sim	sim. years	core-hours	GB/sim yr	output TB	output TB
low	EC60to30	235	60	1024	32.5	40	50	2000	0.13	24	24	6.0
variable	CUSP8	750	80	1024	2.0	20	50	1000	1.02	72	36	9.0
high	RRS18to6	3700	80	1024	0.4	5	25	125	0.67	600	37.5	9.4
idealized		varies							0.18			
total per year									2.00		98	24.4
total for two years									4.00			48.8

Figure 4: Simulation plan. Measured simulated years per day (SYPD) are for stand-alone ocean model without i/o. Total estimate SYPD is doubled to include i/o, sea ice, and coupled climate components.

projects and parameter studies, and for the convenience of working with post-docs and students at LANL.

Requested Resources: Extensive verification and validation is required for changes to fundamental algorithms, such as those proposed here to improve the efficiency of time-stepping. Model domains progress in complexity from simple idealized conditions that test particular physical behavior to full global domains that span from low to high resolution. This ensures that developers can verify against exact solutions and produce convergence tests in simple domains, but still validate against observational data at a range of global resolutions.

The proposed simulations in Figure 4 will be sufficient for these purposes. Throughput totals to 2M CPU hours per year, and is computed based on 1024 cores on grizzly, tested with the MPAS-Ocean code base in June of 2019. The table includes the estimated number and length of simulations for each domain, to give sufficient compute time for thorough validation and parameter sensitivity studies. Simulation output will be kept temporarily on the scratch space, and a subset will be archived for longer-term use.

Conclusion: We propose a total two-year allocation of 4M core-hrs on Badger and Grizzly for the project, allocated as 2M core-hrs each year. Campaign Storage is needed for input data sets totaling 1 TB. Archive Storage needs total 50 TB, with a maximum of 50 TB storage needed on the scratch system at any time.

References

- [1] D. Engwirda. JIGSAW-GEO (1.0): locally orthogonal staggered unstructured grid generation for general circulation modelling on the sphere. *Geoscientific Model Development*, 10(6):2117–2140, June 2017.
- [2] J.-C. Golaz, P. M. Caldwell, L. P. V. Roedel, M. R. Petersen, Q. Tang, and J. D. Wolfe. The DOE E3sm Coupled Model Version 1: Overview and Evaluation at Standard Resolution. *Journal of Advances in Modeling Earth Systems*, 0(0), 2019.
- [3] T.-T.-P. Hoang, W. Leng, L. Ju, Z. Wang, and K. Pieper. Conservative explicit local time-stepping schemes for the shallow water equations. *Journal of Computational Physics*, 382:152 – 176, 2019.
- [4] M. R. Petersen, X. S. Asay-Davis, A. S. Berres, Q. Chen, N. Feige, M. J. Hoffman, D. W. Jacobsen, P. W. Jones, M. E. Maltrud, S. F. Price, T. D. Ringler, G. J. Streletz, A. K. Turner, L. P. V. Roedel, M. Veneziani, J. D. Wolfe, P. J. Wolfram, and J. L. Woodring. An Evaluation of the Ocean and Sea Ice Climate of E3sm Using MPAS and Interannual CORE-II Forcing. *Journal of Advances in Modeling Earth Systems*, 11(5):1438–1458, 2019.
- [5] M. R. Petersen, D. W. Jacobsen, T. D. Ringler, M. W. Hecht, and M. E. Maltrud. Evaluation of the arbitrary Lagrangian–Eulerian vertical coordinate method in the MPAS-Ocean model. *Ocean Modelling*, 86:93–113, Feb. 2015.
- [6] S. M. Reckinger, M. R. Petersen, and S. J. Reckinger. A study of overflow simulations using MPAS-Ocean: Vertical grids, resolution, and viscosity. *Ocean Modelling*, 96:291–313, Dec. 2015.

- [7] T. Ringler, M. Petersen, R. L. Higdon, D. Jacobsen, P. W. Jones, and M. Maltrud. A multi-resolution approach to global ocean modeling. *Ocean Modelling*, 69:211–232, Sept. 2013.
- [8] T. D. Ringler, J. Thuburn, J. B. Klemp, and W. C. Skamarock. A unified approach to energy conservation and potential vorticity dynamics for arbitrarily-structured C-grids. *Journal of Computational Physics*, 229(9):3065–3090, May 2010.
- [9] P. J. Wolfram, T. D. Ringler, M. E. Maltrud, D. W. Jacobsen, and M. R. Petersen. Diagnosing Isopycnal Diffusivity in an Eddying, Idealized Midlatitude Ocean Basin via Lagrangian, in Situ, Global, High-Performance Particle Tracking (LIGHT). *Journal of Physical Oceanography*, 45:2114–2133, Aug. 2015.